

Amendments to the Specification

Please insert the following paragraph beginning at page 1, between lines 2 and 4:

This application is a continuation application of Serial No. 09/453,216, filed December 3, 1999.

Please replace the paragraph beginning on page 1, line 6 with the following rewritten paragraph:

This invention relates to disk array devices and more specifically, to a disk array device in which multiple disks (typically, magnetic disks or optical disks) construct a disk array capable of storing a large volume of data, transferring data at high speed, and further providing higher system reliability.

Please replace the paragraph beginning on page 2, line 7 with the following rewritten paragraph:

Described next is the procedure of creating redundant data with reference to FIGS. 70a and 70b. Data to be written arrives at the controller 6901 by a unit of a predetermined size (2048 bytes, in this description). Here, as shown in FIG. 70a, currently-arrived data is called D-1. The data D-1 is divided into four by the controller 6901, and thereby four data blocks D-A1, D-B1, D-C1, and D-D1 are created. Each data block has a data length of 512 bytes.

Please replace the paragraph beginning on page 2, line 15 with the following rewritten paragraph:

The controller 6901 then creates redundant data D-P1 using the data blocks D-A1, D-B1, D-C1, and D-D1 by executing a calculation given by:

$$D-P1i = D-A1i \text{ xor } D-B1i \text{ xor } D-C1i \text{ xor } D-D1i \dots(1).$$

Please replace the paragraph beginning on page 4, line 1 with the following rewritten paragraph:

There is a possibility that a failure or fault may occur in any disk drives. Here, assuming that the disk drive 6902C has failed and the host device has sent a read request for the data D-1. In this case, the controller 6901 first tries to read the data blocks D-A1, D-B1, D-C1, and D-D1 from the disk drives 6902A, 6902B, 6902C, and 6902D, respectively. However, since the disk drive 6902C is eventually failed, the data block D-C1 is not read therefrom. Assume herein, however, that the data blocks D-A1, D-B1, and D-D1 are read from the disk drives 6902A, 6902B, and 6902D normally. When recognizing that the data block D-C1 cannot be read, the controller 6901 reads the redundant data D-P1 from the disk drive 6902P.

Please replace the paragraph beginning on page 4, line 13 with the following rewritten paragraph:

The controller 6901 then recovers the data block D-C1 by executing a calculation given by the following equation (2) using the data blocks D-A1, D-B1, and D-D1 and the redundant data D-P1.

$$D-C1_i = D-A1_i \text{ xor } D-B1_i \text{ xor } D-D1_i \text{ xor } D-P1_i \dots(2).$$

Please replace the paragraph beginning on page 5, line 10 with the following rewritten paragraph:

In recent years, the RAID architecture, as an example of a disk array, is often implemented also in video servers which provide video upon a user's request. In video servers, data to be stored in the disk drives 6902A to 6902D of the RAID includes two types: video data and computer data (typically, video title and total playing time). Since video data and computer data have different characteristics, requirements of the RAID system are different in reading video data and computer data.

Please replace the paragraph beginning on page 7, line 5 with the following rewritten paragraph:

As shown in FIG. 71a, the first RAID determines that the data-drive D has failed when the data-drive D does not start reading after the lapse of the predetermined time from the start of a fourth reading (data-drive B). To recover the data block of the data-drive D, the first RAID operates calculation of parity. In general disk drives, however, the time from a start to an end of reading is not constant. Some disks may complete reading in a short period of time, while others may take a long time to complete reading after several failures. Therefore, in the first RAID, as shown in FIG. 71b, even though the parity-drive P starts reading earlier than the data-drive B which starts reading fourth, the data-drive B may complete its reading earlier than the parity-drive P. In this case, even after the lapse of the predetermined time after the data-drive B starts reading, the redundant data has not been read from the parity-drive P. Therefore, the first RAID cannot recover the data-block of the data-drive D. As a result, transmission of the data composing the video data being read is delayed, and the video being replayed at the host device might be interrupted.

Please replace the paragraph beginning on page 7, line 24 with the following rewritten paragraph:

A second RAID disclosed in Japanese Patent Laying-Open No. 9-69027 is now described. The second RAID also includes a plurality of data-drives for storing data, and a parity-drive for storing redundant data created from the data. The second RAID does not read the redundant data from the parity-drive under normal conditions. That is, when a read request arrives, the second RAID tries to read the data blocks from the plurality of data-drives. The second RAID previously stores time (hereinafter referred to as predetermined time) by which the plurality of data-drives have to have completed reading. In some cases, the second RAID detects the data-drive which has not completed reading after the lapse of the predetermined time from the time of transmission of a read request to each data-drive. In this case, the second RAID reads the redundant data from the parity-drive to recover the data block which has not yet been completely read.

Please replace the paragraph beginning on page 9, line 3 with the following rewritten paragraph:

In most cases, in the disk drive where reading of the data block is delayed, read requests subsequent to the read request currently being processed wait for a read operation. Therefore, when the disk drive fails to read the data block and retries reading of the data block, processing of the subsequent read requests is delayed. As evident from above, in the conventional disk array device including the above first and second RAID, a read failure may affect subsequent reading.

Please replace the paragraph beginning on page 10, line 12 with the following rewritten paragraph:

However, when another failure occurs in another disk drive of the same parity group while the defective disk drive is left as it is, reconstruction cannot be executed. Therefore, reconstruction is required to be executed as early as possible. An example of such reconstruction is disclosed in Japanese Patent Laying-Open No. 5-127839. A disk array device disclosed in this publication (hereinafter referred to as first disk array device) includes a disk array composing a plurality of disk drives, and a disk controller for controlling the disk array. The disk controller monitors states of operation of the disk array. When reconstruction is required, the disk controller selects and executes one of three types of reconstruction methods according to the state of operation of the disk array. In one method, reconstruction occurs during idle time of the array. In a second method reconstruction is interleaved between current data area accessing operations of the array at a rate which is inversely proportional to an activity level of the array. In a third method, the data are reconstructed when a data area being accessed is a data area needing reconstruction.

Please replace the paragraph beginning on page 12, line 2 with the following rewritten paragraph:

Further, the conventional disk array device may execute reassignment. The structure of a disk array device of executing reassigning is similar to that shown in FIG. 69. Reassignment processing is now described in detail. Each disk drive composing a disk array has recording areas in which a defect may occur due to various reasons. Since the disk drive cannot read/write a data block or

redundant data from/in a defective area, an alternate recording area is reassigned to the defective recording area. In the alternate recording area, the data block or redundant data stored in the defective recording area or to be written in the defective area is stored. Two types of such reassignment have been known.

Please replace the paragraph beginning on page 12, line 14 with the following rewritten paragraph:

One reassignment is so-called auto-reassign executed by each disk drive composing the disk array. Each disk drive previously reserves part of its recording areas as alternate areas. When the data block or redundant data cannot be read/written from/in the recording area specified by the controller, the disk drive assumes that the specified area is defective. When detecting the defective area, the disk drive selects one of the reserved alternate areas, and assigns the selected alternate area to the detected defective area.

Please replace the paragraph beginning on page 13, line 7 with the following rewritten paragraph:

In some recording areas, reading or writing may be eventually successful if the disk drive repeats access to these recording areas (that is, if the disk drive takes much time to access thereto). In the above two types of reassignment, however, the alternate area cannot be assigned to the recording area to which the disk drive takes much time to access because reading/writing will eventually succeed even though much time is required. When the data block composing the video data is stored in such a recording area, however, it takes much time to read the data block. As a result, video being replayed at the host device may be interrupted.

Please replace the paragraphs beginning on page 15, line 2 and ending on page 15, line 7 with the following rewritten paragraphs:

According to a second aspect, in the first aspect, when (m-1) of the disk drives complete reading, the control part:

determines that reading being executed in one remaining disk drive is no longer necessary; and

issues a read termination command to the remaining disk drive.

Please replace the paragraph beginning on page 15, line 8 with the following rewritten paragraph:

As described above, in the second aspect, also when reading of one disk drive takes too much time, this reading is terminated. Therefore, it is possible to provide the disk array device in which, if reading of one disk drive is delayed, this delay does not affect other reading.

Please replace the paragraphs beginning on page 15, line 13 and ending on page 15, line 24 with the following rewritten paragraphs:

According to a third aspect, in the first aspect, when detecting that two or more of the disk drives cannot complete reading, the control part:

determines that reading being executed in other disk drives is no longer necessary; and
issues a read termination command to the determined disk drive.

Please replace the paragraphs beginning on page 15, line 25 and ending on page 16, line 6 with the following rewritten paragraphs:

According to a fourth aspect, in the first aspect, when the (m-1) the disk drives complete reading, the control part:

determines that reading not yet being executed in one remaining disk drive is no longer necessary; and
issues a read termination command to the remaining disk drive.

Please replace the paragraphs beginning on page 16, line 10 and ending on page 17, line 13 with the following rewritten paragraphs:

A fifth aspect of the present invention is directed to a disk array device executing read operation for reading data recorded therein in response to a first read request from a host device, the disk array device with data blocks generated by dividing the data and redundant data generated from the data blocks recorded therein, comprising:

m disk drives across which the data blocks and the redundant data are distributed;
a parity calculation part operating calculation of parity from (m-2) of the data blocks and the redundant data to recover one remaining data block; and
a control part controlling the read operation;
the control part:
issuing second read requests to read the data blocks and the redundant data from the m disk drives in response to the first read request sent thereto;
when (m-1) of the disk drives complete reading, detecting whether a set of the data blocks and the redundant data has been read from the (m-1) disk drives;
when detecting that the set of the data blocks and the redundant data has been read, issuing a recovery instruction to the parity calculation part to recover the data block not read from the one remaining disk drive after waiting for a predetermined time period from a time of detection; and
when the one remaining data block is recovered by the calculation of parity in the parity calculation part, executing operation for transmitting the data to the host device, wherein the predetermined time period is selected so as to ensure data transmission to the host device without delay.

Please replace the paragraphs beginning on page 17, line 21 and ending on page 17, line 25 with the following rewritten paragraph:

According to a sixth aspect, in the fifth aspect, when detecting that the set of the data blocks and the redundant data has not been read, the control part transmits the data to the host device without waiting for the predetermined time period from the a time of detecting.

Please replace the paragraphs beginning on page 18, line 6 and ending on page 18, line 9 with the following rewritten paragraph:

According to a seventh aspect, in the fifth aspect, the predetermined time period is selected based on a start of reading in each of the disk drives and a probability of completing the reading.

Please replace the paragraph beginning on page 18, line 13 and ending on page 19, line 16 with the following rewritten paragraphs:

An eighth aspect of the present invention is directed to a disk array device executing read operation for reading data recorded therein in response to a first read request from a host device, the disk array device with data blocks generated by dividing the data and redundant data generated from the data blocks recorded therein, comprising:

m disk drives across which the data blocks and the redundant data are distributed;

a parity calculation part operating calculation of parity from (m-2) of the data blocks and the redundant data to recover one remaining data block; and

a control part controlling the read operation; the control part:

issuing second read requests to read the data blocks and the redundant data from the m disk drives in response to the first read request sent thereto;

when (m-1) of the disk drives complete reading, detecting whether a set of the data blocks and the redundant data has been read from the (m-1) disk drives;

when detecting that the set of the data blocks and the redundant data has been read, issuing a recovery instruction to the parity calculation part to recover the data block not read from the one remaining disk drive after waiting for a predetermined time period from a time of detection; and

when the one remaining block is recovered by the calculation of parity in the parity calculation part, executing operation for transmitting the data to the host device, wherein the recovery instruction is issued while the parity calculation part is not operating calculation of parity.

Please replace the paragraphs beginning on page 19, line 21 and ending on page 20, line 5 with the following rewritten paragraphs:

According to a ninth aspect, in the eighth aspect, the disk array device further comprises:

a table including a time period during which the parity calculation part can operate calculation of parity, wherein the control part further issues the recovery instruction when the parity calculation part does not operate calculation of parity by referring to the time period included in the table.

Please replace the paragraphs beginning on page 20, line 6 and ending on page 21, line 3 with the following rewritten paragraphs:

A tenth aspect of the present invention is directed to a disk array device executing read operation for reading data recorded therein in response to a first read request from a host device, the disk array device with data blocks generated by dividing the data and redundant data generated from the data blocks recorded therein, comprising:

- m disk drives across which the data blocks and the redundant data are distributed;
- a parity calculation part operating calculation of parity from (m-2) of the data blocks and the redundant data to recover one remaining data block; and
- a control part controlling the read operation, the control part:
 - in response to the first read request received thereto, determining whether (m-1) of the disk drives have previously failed to read each data block or not;
 - when determining that the (m-1) disk drives have not previously failed to read each of the data block, issuing second read requests to the (m-1) disk drives to read only each the data blocks; and
 - when the data blocks are read from the (m-1) disk drives, executing operation for transmitting the data to the host device.

Please replace the paragraphs beginning on page 21, line 9 and ending on page 21, line 25 with the following rewritten paragraphs:

According to an eleventh aspect, in the tenth aspect, the control part:

- when determining that the (m-1) disk drives have previously failed to read each the data block, issues second read requests to the m disk drives to read (m-1) of the data blocks and the redundant data;
- when the (m-1) disk drives complete reading, detects whether a set of the data blocks and the redundant data has been read from the (m-1) disk drives or not;
- when detecting that the set of the data blocks and the redundant data has been read, issues a recovery instruction to the parity calculation part to recover the data block not read from one remaining disk drive; and

when the one remaining data block is recovered by the calculation of parity in the parity calculation part, executes operation for executing operation for transmitting the data to the host device.

Please replace the paragraphs beginning on page 22, line 4 and ending on page 22, line 13 with the following rewritten paragraphs:

According to a twelfth aspect, in the eleventh aspect, the disk array device further comprises:
a table registering therein recording areas of the data blocks which have previously been failed to be read by the disk drives, wherein the control part determines whether to issue the second read requests to the (m-1) disk drives or to the m disk drives.

Please replace the paragraphs beginning on page 22, line 14 and ending on page 22, line 23 with the following rewritten paragraphs:

According to a thirteenth aspect, in the twelfth aspect, the disk array device further comprises:
a reassignment part, when a defect occurs in a recording area of the data block or redundant data in the m disk drives, executing reassign processing for assigning an alternate recording area to the defective recording area, wherein when the reassignment part assigns the alternate recording area to the defective recording area of the data block registered in the table by the reassignment part, the control part deletes the defective recording area of the data block from the table.

Please replace the paragraphs beginning on page 23, line 6 and ending on page 24, line 9 with the following rewritten paragraphs:

According to a fourteenth aspect, in the thirteenth aspect, the disk array device further comprises:

a first table storage part storing a first table in which an address of the alternate recording area previously reserved in each of the m disk drives can be registered as alternate recording area information; and

a second table storage part storing a second table in which address information of the alternate recording area assigned to the defective recording area can be registered, wherein the reassignment part:

when the second read requests are transmitted from the control part to the m disk drives, measures a delay time in each of the disk drives;

determines whether each of the recording area of the data blocks or the redundant data to be read by each second read request is defective or not based on the measured delay time;

when determining that the recording area is defective, assigns the alternate recording area to the defective recording area based on the alternate recording area information registered in the first table of the first table storage part; and

registers the address information of the assigned alternate recording area in the second table of the second table storage part;

the control part issues the second read requests based on the address information registered in the second table of the second table storage part; and

the delay time is a time period calculated from a predetermined process start time.

Please replace the paragraphs beginning on page 24, line 24 and ending on page 26, line 1 with the following rewritten paragraphs:

According to a sixteenth aspect, in the fifteenth aspect, the disk array device further comprises:

a first table storage part storing a first table in which an address of the alternate recording area previously reserved in each of the m disk drives can be registered as alternate recording area information; and

a second table storage part storing a second table in which address information of the alternate recording area assigned to the defective recording area can be registered, wherein the reassignment part:

when the second read requests are transmitted from the control part to the m disk drives, measures a delay time in each of the disk drives;

determines whether each of the recording areas of the data blocks or the redundant data to be read by each second read request is defective or not based on the measured delay time;

when determining that the recording area is defective, assigns the alternate recording area to the defective recording area based on the alternate recording area information registered in the first table of the first table storage part; and

registers the address information of the assigned alternate recording area in the second table of the second table storage part;

the control part issues the second read requests based on the address information registered in the second table of the second table storage part; and

the delay time is a time period calculated from a predetermined process start time.

Please replace the paragraphs beginning on page 25, line 2 and ending on page 25, line 6 with the following rewritten paragraph:

According to a seventeenth aspect, in the sixteenth aspect, the reassignment part assigns the alternate recording area to the defective recording area only when determining successively a predetermined number of times that the recording area is defective.

Please replace the paragraphs beginning on page 26, line 16 and ending on page 26, line 19 with the following rewritten paragraph:

According to an eighteenth aspect, in the sixteenth aspect, the predetermined process start time is a time when each of the second read requests is transmitted to each of the m disk drives.

Please replace the paragraphs beginning on page 26, line 20 and ending on page 26, line 22 with the following rewritten paragraph:

According to a nineteenth aspect, in the sixteenth aspect, the predetermined process start time is a time when the m disk drives start reading based on the second read requests.

Please replace the paragraphs beginning on page 26, line 25 and ending on page 27, line 17 with the following rewritten paragraphs:

A twentieth aspect of the present invention is directed to a data input/output method used for a disk array device comprising a disk array constructed of recording mediums for recording redundant data and an array controller for controlling the disk array according to an access request transmitted from a host device, the method comprising:

generating, by the array controller, a read or write request to the disk array with predetermined priority based on the received access request;

enqueueing, by the array controller, the generated read or write request to a queue included therein according to the predetermined priority;

selecting, by the array controller, the read or write request to be processed by the disk array from among the read or write requests enqueueued to the queue according to the predetermined priority; and

processing, by the disk array, the selected read or write request.

Please replace the paragraphs beginning on page 28, line 8 and ending on page 28, line 13 with the following rewritten paragraphs:

According to a twenty-first aspect, in the twentieth aspect, the array controller includes queues therein corresponding to the priority; and

the generated read request or write request is enqueueued to the queue corresponding to the predetermined priority.

Please replace the paragraphs beginning on page 28, line 19 and ending on page 29, line 4 with the following rewritten paragraphs:

According to a twenty-second aspect, in the twentieth aspect, the array controller includes queues therein corresponding to the predetermined priority for each of the recording mediums, the array controller generates the read or write request with the predetermined priority for each of the recording mediums 25 based on the received access request; and

the array controller enqueues the read or write request generated for each of the recording mediums to the queue in the corresponding recording medium according to the predetermined priority.

Please replace the paragraph beginning on page 29, line 11 and ending on page 29, line 14 with the following rewritten paragraph:

According to a twenty-third aspect, in the twentieth aspect, the predetermined priority is set based on whether processing in the disk array is executed in real time or not.

Please replace the paragraphs beginning on page 29, line 20 and ending on page 29, line 25 with the following rewritten paragraph:

According to a twenty-fourth aspect, in the twentieth aspect, when an I/O interface is between the information recording device and the host device conforms to SCSI, the predetermined priority is previously set in a LUN or LBA field of the access request.

Please replace the paragraphs beginning on page 30, line 6 and ending on page 30, line 21 with the following rewritten paragraph:

A twenty-fifth aspect of the present invention is directed to a disk array device including a disk array constructed of recording mediums for recording redundant data and controlling the disk array according to an access request transmitted from a host device, comprising:

- a control part generating a read or write request to the disk array with predetermined priority based on the received access request;

- a queue managing part enqueueing the read request or write request generated by the control part to a queue included therein according to the predetermined priority; and

- a selection part selecting the read or write request to be processed by the disk array from among the read or write requests enqueued to the queue, wherein the disk array processes the read request or write request selected by the selection part.

Please replace the paragraphs beginning on page 31, line 12 and ending on page 31, line 18 with the following rewritten paragraph:

According to a twenty-sixth aspect, in the twenty-fifth aspect, the queue managing part includes queues therein corresponding to the priority, and the read or write request generated by the control part is enqueued to the queue corresponding to the predetermined priority.

Please replace the paragraphs beginning on page 31, line 24 and ending on page 32, line 10 with the following rewritten paragraphs:

According to a twenty-seventh aspect, in the twenty-fifth aspect, the queue managing part includes queues therein corresponding to the predetermined priority for each of the recording mediums;

the queue managing part generates the read or write request with the predetermined priority for each of the recording mediums based on the received access request; and

the queue managing part enqueues the read or write request generated for each of the recording mediums to the queue in the corresponding recording medium according to the predetermined priority.

Please replace the paragraphs beginning on page 32, line 17 and ending on page 33, line 13 with the following rewritten paragraphs:

A twenty-eighth aspect of the present invention is directed to, in an information recording device comprising a disk array constructed of recording mediums for recording redundant data and an array controller for controlling the disk array according to an access request transmitted from a host device, a data reconstruction method for recovering data recorded on a faulty recording medium in the disk array and reconstructing the data, the method comprising:

generating, by the array controller, a read or write request required for data reconstruction to the disk array with predetermined priority;

enqueueing, by the array controller, the generated read or write request to a queue included therein according to the predetermined priority;

selecting, by the array controller, the read or write request to be processed from among the read or write requests enqueued to the queue according to the predetermined priority;
processing, by the disk array, the selected read or write request; and
executing, by the array controller, data reconstruction based on processing results of the read or write request by the disk array.

Please replace the paragraphs beginning on page 34, line 1 and ending on page 34, line 10 with the following rewritten paragraphs:

According to a twenty-ninth aspect, in the twenty-eighth aspect, the array controller includes queues therein corresponding to the predetermined priority for each of the recording mediums;
the array controller generates the read or write request required for data reconstruction with the predetermined priority for each recording medium; and
the array controller enqueues the generated read or write request to the queue in the corresponding recording medium according to the predetermined priority.

Please replace the paragraphs beginning on page 34, line 19 and ending on page 34, line 23 with the following rewritten paragraph:

According to a thirtieth aspect, in the twenty-eighth aspect, the read and write requests generated by the array controller are given lower priority to be processed in the disk array.

Please replace the paragraphs beginning on page 35, line 2 and ending on page 35, line 6 with the following rewritten paragraph:

According to a thirty-first aspect, in the twenty-eighth aspect, the read and write requests generated by the array controller are given higher priority to be processed in the disk array.

Please replace the paragraphs beginning on page 35, line 10 and ending on page 36, line 15 with the following rewritten paragraphs:

A thirty-second aspect of the present invention is directed to a data input/output method used in an information recording device comprising a disk array constructed of recording mediums for

recording redundant data and an array controller for controlling the disk array according to an access request transmitted from a host device, recovering the data recorded on the recording medium which has a failure in the disk array, and reconstructing the data in a spare recording medium;

when the access request for data to be reconstructed in the spare recording medium is transmitted from the host device to the information storage device, the method comprises:

the array controller;

reading data for recovery required for recovering the data recorded in the failed recording medium from the disk array;

recovering data recorded in the failed recording medium by executing a predetermined calculation with the data for recover read from the disk array;

generating a write request with predetermined priority to write the recovered data in the spare recording medium;

enqueuing the generated write request to a queue therein according to the predetermined priority; and

selecting the generated write request as the write request to be processed by the disk array according to the predetermined priority; and

the disk array:

processing the write request selected by the array controller, and writing the recovered data in the spare recording medium, wherein the write request is given relatively lower priority.

Please replace the paragraphs beginning on page 36, line 23 and ending on page 37, line 15 with the following rewritten paragraphs:

A thirty-third aspect of the present invention is directed to a disk array device which reassigns an alternate recording area to a defective recording area of data, comprising:

a read/write control part for specifying a recording area of data, and producing an I/O request to request read or write operation;

a disk drive, when receiving the I/O request transmitted from the read/write control part, accessing to the recording area specified by the I/O request to read or write the data; and

a reassignment part when receiving the I/O request transmitted from the read/write control part, calculating an elapsed time from a predetermined process start time as a delay time and determining whether the recording area specified by the I/O request is defective or not based on the delay time, wherein when determining that the recording area of the data is defective, the reassignment part instructs the disk drive to assign the alternate recording area to the defective recording area.

Please replace the paragraphs beginning on page 38, line 5 and ending on page 38, line 10 with the following rewritten paragraph:

According to a thirty-fourth aspect, in the thirty-third aspect, the reassignment part assigns the alternate recording area to the defective recording area only when determining successively a predetermined number of times that the recording area is defective.

Please replace the paragraphs beginning on page 38, line 19 and ending on page 38, line 22 with the following rewritten paragraph:

According to a thirty-fifth aspect, in the thirty-third aspect, the predetermined process start time is a time when the I/O request is transmitted from the read/write control part.

Please replace the paragraphs beginning on page 38, line 23 and ending on page 39, line 2 with the following rewritten paragraph:

According to a thirty-sixth aspect, in the thirty-third aspect, the predetermined process start time is a time when the I/O request transmitted from the read/write control part is started to be processed in the disk drive.

Please replace the paragraphs beginning on page 39, line 8 and ending on page 39, line 12 with the following rewritten paragraph:

According to a thirty-seventh aspect, in the thirty-third aspect, the reassignment part further instructs the disk drive to terminate the read or write operation requested by the I/O request when the recording area of the data is defective.

Please replace the paragraphs beginning on page 39, line 20 and ending on page 40, line 11 with the following rewritten paragraphs:

A thirty-eighth aspect of the present invention is directed to a disk array device which reassigns an alternate recording area to a defective recording area of data, comprising:

a read/write control part specifying a recording area of the data, and producing an I/O request to request read or write operation;

a disk drive, when receiving the I/O request from the read/write control part, accessing to the recording area specified by the I/O request to read or write the data; and

a reassignment part, when the recording area specified by the I/O request from the read/write control part is defective, instructing the disk drive to reassign the alternate recording area to the defective recording area, wherein when instructed to reassign by the reassignment part, the disk drive assigns a recording area in which time required for the read or write operation is within a predetermined range, as the alternate recording area.

Please replace the paragraphs beginning on page 40, line 18 and ending on page 40, line 21 with the following rewritten paragraph:

According to a thirty-ninth aspect, in the thirty-eighth aspect, the predetermined range is selected based on overhead in the disk array device.

Please replace the paragraphs beginning on page 41, line 1 and ending on page 41, line 5 with the following rewritten paragraph:

According to a fortieth aspect, in the thirty-eighth aspect, when part or all of the recording areas of the data are defective, the reassignment part assumes that the whole recording areas are defective.

Please replace the paragraph beginning on page 41, line 6 with the following rewritten paragraph:

In the fortieth aspect, in the disk array device, the alternate recording area is assigned not by a fixed-block unit, which is a managing unit in the disk drive. Therefore, the disk array device can prevent data fragmentation, suppressing occurrence of a large delay in response more.

Please replace the paragraphs beginning on page 41, line 11 and ending on page 41, line 20 with the following rewritten paragraphs:

According to a forty-first aspect, in the thirty-eighth aspect, the reassignment part transmits a reassign block specifying a logical address block of the defective recording area to the disk drive for reassignment; and

the disk drive assigns a physical address with which the time required for read or write operation is within the predetermined range to a logical address specified by the reassign block transmitted from the reassignment part as the alternate recording area.

Please replace the paragraphs beginning on page 42, line 3 and ending on page 42, line 11 with the following rewritten paragraphs:

According to a forty-second aspect, in the thirty-eighth aspect, when the read/write control part requests the disk drive to read the data, and the recording area of the data is defective, the data recorded in the defective recording area is recovered based on predetermined parity and other data; and

the read/write control part specifies the assigned alternate recording area, and requests the disk drive to write the recovered data.

Please replace the paragraphs beginning on page 42, line 12 and ending on page 42, line 18 with the following rewritten paragraph:

According to a forty-third aspect, in the thirty-eighth aspect, when the read/write control part requests the disk drive to write data and the recording area of the data is defective, the read/write

control part specifies the assigned alternate recording area, and the requests again the disk drive to write the data.

Please replace the paragraphs beginning on page 43, line 3 and ending on page 43, line 15 with the following rewritten paragraphs:

A forty-fourth aspect of the present invention is directed to a reassignment method of assigning an alternate area to a defective recording area of data, comprising:

transmitting an I/O request for requesting the disk drive to read or write operation by specifying a recording area of the data according to a request from outside; and

when the I/O request is transmitted in the transmission step, calculating an elapsed time from a predetermined time as a delay time and determining whether the recording area specified by the I/O request is defective or not based on the delay time, wherein when the recording area is defective in the determination step, the disk drive is instructed to assign the alternate recording area to the defective recording area.

Please replace the paragraphs beginning on page 43, line 16 and ending on page 44, line 4 with the following rewritten paragraphs:

A forty-fifth aspect of the present invention is directed to a reassignment method of assigning an alternate recording area to a defective recording area of data, comprising:

transmitting an I/O request for requesting the disk drive to read or write operation by specifying a recording area of the data according to a request from outside; and

when the recording area specified by the I/O request transmitted in the transmission step is defective, instructing the disk drive to assign the alternate recording area to the defective recording area, wherein in the instructing step, the disk drive is instructed to assign the recording area with which time required for read or write operation is within a predetermined range as the alternate recording area.

Please replace the paragraphs beginning on page 44, line 5 and ending on page 45, line 7 with the following rewritten paragraphs:

A forty-sixth aspect of the present invention is directed to a disk array device which assigns an alternate recording area to a defective recording area of data, comprising:

a read/write control part for transmitting an I/O request for requesting read or write operation by specifying a recording area of the data according to a request from outside;

a disk drive, when receiving the I/O request from the read/write control part, accessing to the recording area specified by the I/O request and reading or writing the data;

a reassignment part, when receiving the I/O request from the read/write control part, calculating an elapsed time from a predetermined process start time as a delay time, and determining whether the recording area specified by the I/O request is defective or not based on the delay time;

a first storage part storing an address of the alternate recording area previously reserved in the disk drive as alternate recording area information; and

a second storage part storing address information of the alternate recording area assigned to the defective recording area, wherein when determining that the specified recording area is defective, the reassignment part assigns the alternate recording area to the defective recording area based on the alternate recording area information stored in the first storage part, and stores the address information on the assigned alternate recording area in the second storage part, and the read/write control part generates the I/O request based on the address information stored in the second storage part.

Please replace the paragraphs beginning on page 45, line 17 and ending on page 45, line 22 with the following rewritten paragraph:

According to a forty-seventh aspect, in the forty-sixth aspect, the reassignment part assigns the alternate recording area to the defective recording area only when determining successively a predetermined number of times that the recording area is defective.

Please replace the paragraphs beginning on page 45, line 23 and ending on page 46, line 1 with the following rewritten paragraph:

According to a forty-eighth aspect, in the forty-sixth aspect, the predetermined process start time is a time when the I/O request is transmitted from the read/write control part.

Please replace the paragraphs beginning on page 46, line 2 and ending on page 46, line 6 with the following rewritten paragraph:

According to a forty-ninth aspect, in the forty-sixth aspect, the predetermined process start time is a time when the I/O request transmitted from the read/write control part is started to be processed in the disk drive.

Please replace the paragraphs beginning on page 46, line 7 and ending on page 46, line 10 with the following rewritten paragraph:

According to a fiftieth aspect, in the forty-sixth aspect, the reassignment part further instructs the disk drive to terminate the read or write operation requested by the I/O request when detecting that the recording area of the data is defective.

Please replace the paragraphs beginning on page 46, line 11 and ending on page 46, line 15 with the following rewritten paragraph:

According to a fifty-first aspect, in the forty-sixth aspect, the first storage part stores a recording area with which overhead in the disk drive is within a predetermined range as the alternate recording area.

Please replace the paragraphs beginning on page 47, line 2 and ending on page 47, line 6 with the following rewritten paragraph:

According to a fifty-second aspect, in the fifty-first aspect, the first storage part further stores the alternate recording area by a unit of a size of the data requested by the I/O request.

Please replace the paragraphs beginning on page 47, line 13 and ending on page 47, line 20 with the following rewritten paragraph:

According to a fifty-third aspect, in the fifty-second aspect, whether the overhead is within the predetermined range or not is determined for the recording areas other than the alternate recording area by the unit, and the reassignment part assigns the alternate area to the recording area in which the overhead is not within the predetermined range.

Please replace the paragraphs beginning on page 48, line 5 and ending on page 48, line 8 with the following rewritten paragraph:

According to a fifty-fourth aspect, in the forty-sixth aspect, the address information stored in the second storage part is recorded in the disk drive.

Please replace the paragraphs beginning on page 48, line 25 and ending on page 49, line 9 with the following rewritten paragraphs:

According to a fifty-sixth aspect, in the forty-sixth aspect, the disk array device further comprises:

a plurality of disk drives including data recording disks device and a spare disk drive; and
a count part counting a used amount or remaining amount of alternate recording area, wherein the reassignment part determines whether to copy the data recorded in the data recording disk drives to the spare disk drive based on a count value in a count part, thereby allowing the spare disk drive to be used instead of the data recording disk drive.

Please replace the paragraphs beginning on page 49, line 17 and ending on page 50, line 10 with the following rewritten paragraphs:

A fifty-seventh aspect of the present invention is directed to a reassignment method of assigning an alternate recording area to a defective recording area of data, comprising:

transmitting an I/O request for requesting read or write operation by specifying a recording area of the data; and

when the recording area specified by the I/O request transmitted in the transmission step is defective, assigning the alternate recording area to the defective recording area, wherein in the assign step, when the specified recording area is defective, the alternate recording area is selected for the defective recording area by referring to alternate recording area information for managing an address of the alternate recording area previously reserved in the disk drive, the selected alternate recording area is assigned to the defective recording area, and further address information for managing an address of the assigned alternate recording area is created; and

in the transmission step, the I/O request is generated based on the address information created in the assign step.

Please replace the paragraphs beginning on page 50, line 11 and ending on page 50, line 17 with the following rewritten paragraph:

According to a fifty-eighth aspect, in the fifty-seventh aspect, in the assign step, when the I/O request is transmitted, an elapsed time from a predetermined process start time is calculated as a delay time, and it is determined whether the recording area specified by the I/O request is defective or not based on the delay time.

Please replace the paragraph beginning on page 51, line 10 with the following rewritten paragraph:

FIGS. 6a and 6b are diagrams illustrating a change in reading order in disk drives 5A to 5D and 5P shown in FIG. 1;

Please replace the paragraph beginning on page 58, line 9 with the following rewritten paragraph:

Each storage area of the other buffer memories 3B to 3D, 3P, and 3R is also divided into a plurality of buffer areas. The identifier is also allocated to each buffer area in the same manner as described for the buffer area 3A₁.

Please replace the paragraph beginning on page 58, line 22 with the following rewritten paragraph:

In response to the write request and data, the RAID starts a write operation. Already being described in detail in Background Art section, this write operation is briefly described herein for the first embodiment with reference to FIGS. 3a and 3b. Assume that transmission data D-1 (refer to FIG. 3a) is inputted from the host device through the host interface 1 to the selector 2 of the disk array device. The selector 2 divides the data D-1 into four, generating data blocks D-A1, D-B1, D-C1, and D-D1 of 512 bytes each. The selector 2 transfers the data block D-A1 to the buffer memory 3A, the data block D-B1 to the buffer memory 3B, the data block D-C1 to the buffer memory 3C, and the data block D-D1 to the buffer memory 3D. The buffer memories 3A to 3D store the transferred data blocks D-A1 to D-D1, respectively.

Please replace the paragraph beginning on page 61, line 21 with the following rewritten paragraph:

In response to the first read request, the disk array device starts a read operation that is distinctive of the present embodiment, which is now described in detail with reference to flow charts in FIGS. 4a and 4b.

Please replace the paragraph beginning on page 65, line 15 with the following rewritten paragraph:

Consider next a case where the controller 7 receives the first READ-COMPLETED from the buffer memory 3P. In this case, the redundant data and three data blocks are stored in the disk drive 5P and three of the disk drives, but one data block has not yet been read. The controller 7 therefore determines that the data required by the host device cannot be transmitted until the unread data block is recovered. The controller 7 then advances from step S13 to step S14, producing a recovery instruction to request the parity calculator 6 to operate calculation of parity (step S14).

Please replace the paragraph beginning on page 65, line 25 with the following rewritten paragraph:

In response to the recovery instruction, the parity calculator 6 fetches the redundant data and three data blocks from the buffer memory area $3P_i$ and three buffer memory areas (any of three buffer areas $3A_i$ to $3D_i$) which store these data blocks. The parity calculator 6 operates calculation of parity as described in the Background Art section to recover the unread data block from the redundant data and three data blocks. The recovered data block is stored in a buffer memory area $3R_i$ in the buffer memory 3R. When the calculation of parity ends, the parity calculator 6 issues a recovery-completed signal indicating an end of calculation of parity, and transmits the signal to the controller 7. When receiving the recovery-complete signal (step S15), the controller 7 determines that four data blocks are stored in the buffer memory areas and that the data requested from the host device can be transmitted. The procedure then advances to step S16.

Please replace the paragraph beginning on page 68, line 19 with the following rewritten paragraph:

To highlight the distinctive characteristics of the present disk array device, described next is read operation by a disk array device which does not execute step S12 of FIG. 4b (hereinafter referred to as a no-termination disk array device), with reference to FIG. 5b. FIG. 5b is a schematic diagram showing read timing of the parity groups n and $(n+1)$ in a time axis in the no-termination array disk device. The conditions in FIG. 5b are the same as those in FIG. 5a except that the no-termination disk array device does not execute step S12 of FIG. 4b. The host device requests data reading from the parity group n , and then the parity group $(n+1)$, under the same conditions as described above.

Please replace the paragraph beginning on page 69, line 5 with the following rewritten paragraph:

The controller 7 issues a set of second read requests in the order in which the first read requests arrive to read data from the parity groups n and $(n+1)$. As shown in FIG. 5b, like in FIG. 5a, reading of the data blocks or redundant data starts in the order as the disk drives 5D, 5C, 5A, 5P,

and 5B. The disk drives 5C, 5A, 5P, and 5B have completed reading by the time t_1 , as is the same in the FIG. 5a, while the disk drive 5D continues reading. Without a read-termination command, reading of the disk drive 5D is not forcefully terminated immediately after the time t_1 , ending at a time t_4 long after the time t_1 . Note that the data of the parity group n can be transmitted to the host device at the time t_1 , as shown in FIG. 5a.

Please replace the paragraph beginning on page 71, line 6 with the following rewritten paragraph:

As clear from above, according to the disk array device of the present invention, the whole volume of data read from the disk drives 5A to 5P (so-called disk array) per unit of time increases. Therefore, the present disk array device can continuously transmit data to the host device. As a result, video data being replayed at the host device has a lower tendency to be interrupted.

Please replace the paragraph beginning on page 71, line 23 with the following rewritten paragraph:

Consider that the controller 7 issues second read requests for reading the data block or redundant data to the disk drive of FIG. 6a in the order as the parity groups n , $(n+1)$, $(n+2)$, $(n+3)$, and $(n+4)$. The disk drive of FIG. 6a executes reading so as to shorten a seek distance of a read head without reading in the order in which the second read requests arrive. For example, the disk drive changes the order of reading so that the read head moves linearly from the inner to outer radius of the disk. As a result, the data blocks and redundant data are read in the order as the parity groups n , $(n+2)$, $(n+4)$, $(n+1)$, and $(n+3)$. The disk drive thus can efficiently read more data blocks and redundant data per unit of time.

Please replace the paragraph beginning on page 72, line 18 with the following rewritten paragraph:

First, the controller 7 issues second read requests as in the requested order. Therefore, the second read requests arrive in each of the disk drives 5A to 5D and 5P in the order as the parity groups n , $(n+1)$, $(n+2)$, $(n+3)$, and $(n+4)$. The disk drives 5A to 5D and 5P, however, determine the

order of reading independently, and thus the actual reading order in each disk drive is not necessarily equal to the requested order and may be different from one another. Furthermore, in FIG. 7a, the disk drives 5A, 5B, and 5P have completed reading the data blocks and redundant data of the parity group (n+2) by a time t_7 and the disk drive 5D completes reading the data block of the same parity group at the time t_7 (refer to hatched parts), while the disk drive 5C completes reading the data block of the parity group (n+4) at the time t_7 (refer to a horizontally-lined part). In this case, the controller 7 receives the fourth first READ-COMPLETED for the parity group (n+2) immediately after the time t_7 (step S11 of FIG. 4b). Therefore, a read termination command is sent to the disk drive 5C (step S12), which therefore does not read the data block of the parity group (n+2).

Please replace the paragraph beginning on page 73, line 19 with the following rewritten paragraph:

To highlight the distinctive characteristics of the present 20 disk array device, described next is read operation by a disk array device which does not execute step S12 of FIG. 4b, with reference to FIG. 7b. FIG. 7b is a schematic diagram showing read timing of the parity groups n to (n+4) in a time axis in the disk array device. The conditions in FIG. 7b are the same as those in FIG. 7a except that the disk array device does not execute step S12 of FIG. 4b. The host device requests data reading from the parity groups n, (n+1), (n+2), (n+3) and then (n+4) sequentially in this order under the same conditions as described above.

Please replace the paragraph beginning on page 74, line 23 with the following rewritten paragraph:

As clear from the above, when a data block becomes in a state of being recoverable, the disk array device of the present invention sends a read termination command to the disk drive which has not yet started reading the data block. In response to the read termination command, the disk device will not start unnecessary reading, and starts only necessary reading. Therefore, the present disk array device can quickly transmit the requested data to the host device. In FIG. 7a, four pieces of data of the parity groups n, (n+2), (n+4), and (n+1) can be transmitted to the host device at a time

t_9 . On the other hand, in FIG. 7b, with unnecessary reading by the disk drives 5C and 5D, only three pieces of data n , $(n+2)$, and $(n+4)$ can be transmitted 10 at the time t_9 .

Please replace the paragraph beginning on page 75, line 11 with the following rewritten paragraph:

As clear from above, according to the disk array device of the present invention, the volume of data to be read per unit of time increases, and data can be continuously transmitted to the host device. As a result, video data being replayed at the host device has a lower tendency to be interrupted.

Please replace the paragraph beginning on page 76, line 4 with the following rewritten paragraph:

Described next is a disk array device according to a second embodiment of the present invention. The configuration of the disk array device is the same as that shown in FIG. 1. For a clear understanding of technical effects of the second embodiment, any of the disk drives 5A to 5D and 5P does not execute reading in the arrival order but changes the reading order so as to shorten the seek distance (the distance required for seeking) of the read head as shown in FIG. 6b.

Please replace the paragraph beginning on page 76, line 18 with the following rewritten paragraph:

In response to the first read request, the disk array device starts a read operation that is distinctive of the present embodiment, which is now described in detail with reference to flowcharts in FIGS. 8a and 8b. Since the flowchart in FIG. 8a partially includes the same steps as those in FIG. 4a, the steps in FIG. 8a are provided with the same step numbers as those in FIG. 4a and their description is simplified herein.

Please replace the paragraph beginning on page 81, line 5 with the following rewritten paragraph:

Assume that a time t_{11} equals to $t_{10} + T_{LIMIT}$ and $(t_{PRE} - t_{ISSUE}) > T_{LIMIT}$ is satisfied. At the time t_{11} , the controller 7 fetches the information on the buffer memory areas $3A_1$ to $3D_1$ and $3P_1$ written with the issue time t_{ISSUE} (t_{10}) from the issue time table 71_n (refer to FIG. 9). By the time t_{11} , only the disk drive 5D has completed reading of the data block of the parity group n , and therefore the controller 7 has received only the first READ-COMPLETED specifying the buffer memory area $3D_1$ from the buffer memory 3D. The controller 7 thus recognizes that two or more first READ-COMPLETED's have not arrived by the limit time T_{LIMIT} and that reading of the parity group n in the disk drives 5A to 5C and 5P has not yet been completed. The controller 7 thus specifies the disk drives (in this case, the disk drives 5A to 5C and 5P) which are taking too much time to read the data of the parity group n .

Please replace the paragraph beginning on page 83, line 14 with the following rewritten paragraph:

As described above, the disk array device does not execute the processing shown in FIG. 8b. Therefore, the disk drives 5A and 5P do not terminate the read operation even though they take longer time than the limit time t_{LIMIT} to read the parity group n . Furthermore, it is highly possible that the data blocks of the parity group n stored in the disk drives 5A and 5P may have a failure. Therefore, the disk array device cannot assemble and transmit the data of the parity group n . Here, note that, despite that, the disk drives 5B and 5C start unnecessary reading of the data block of the parity group n .

Please replace the paragraph beginning on page 83, line 24 with the following rewritten paragraph:

As evident from FIGS. 10a and 10b, with execution of the processing of FIG. 8b, upon realizing that data being read cannot be transmitted to the host device, the disk array device of the second embodiment terminates all reading of the parity group. Therefore, in the case of FIG. 10a, the disk drives 5A, 5B, 5C, and 5P can start reading the next parity group earlier than the case of

FIG. 10b, thereby terminating unnecessary reading and quickly starting the next reading. Further, the disk drives 5B and 5C skip reading of the parity group data of which cannot be transmitted to the host device, and start reading of the next parity group. As a result, the disk array device can read a larger volume of data per unit of time, and thus continuously transmit data to the host device, allowing video data being replayed at the host device to have a lower tendency to be interrupted.

Please replace the paragraph beginning on page 84, line 15 with the following rewritten paragraph:

In the previous embodiments, the controller 7 immediately issues a recovery instruction to the parity calculator 6 after three data blocks and the redundant data are stored in the buffer memories. However, the calculation of parity requires a large amount of arithmetic operation, and the larger the number of operations of calculation of parity, the more the disk array device is loaded. In a disk array device of a third embodiment, the controller 7 controls timing of issuing a recovery instruction to reduce the number of operation of calculation of parity.

Please replace the paragraph beginning on page 85, line 6 with the following rewritten paragraph:

The disk array device performs write operations as described in the first embodiment whenever transmission data arrives from the host device. To read data from the disk array device, the host device transmits a first read request specifying storage locations of the data to the disk array device.

Please replace the paragraph beginning on page 85, line 11 with the following rewritten paragraph:

In response to the first read request, the disk array device starts a read operation that is distinctive of the third embodiment, which is now described in detail with reference to flow charts of FIGS. 12a and 12b. Note that since the flow chart of FIG. 12a is equal to that of FIG. 8a, the steps in FIG. 12a are provided with the same step numbers as those in FIG. 8a. Through the execution of the flow chart of FIG. 12a, the controller 7 issues a set of second read requests (requests for reading

a parity group) (steps S1 and S2), and further creates the issue time table 71 for the issued second read requests (step S21).

Please replace the paragraph beginning on page 88, line 12 with the following rewritten paragraph:

When four first READ-COMPLETED's have arrived at the controller 7, the progress of reading in the disk drives 5A to 5D and 5P is as shown in FIG. 13b, for example. In FIG. 13b, the second read requests issued at the time t_{ISSUE} cause each disk drive to start reading. The disk drives 5A, 5B, 5D, and 5P have completed reading by a time $t_{4\text{th}}$.

Please replace the paragraph beginning on page 88, line 24 with the following rewritten paragraph:

Therefore, the controller 7, in step S43, first fetches the time $t_{4\text{th}}$ stored in step S41, the time t_{ISSUE} in the issue time table 71, and the previously-stored completion-expectation value t_0 . Then, $\{t_0 - (t_{4\text{th}} - t_{\text{ISSUE}})\}$ is calculated, resulting in a time margin t_{MARGIN} as shown in a hollow double-headed arrow in FIG. 13b. The controller 7 sets the first timer 72 to the calculated time margin t_{MARGIN} as the timeout value V_{TOI} (step S43 in FIG. 12b). This activates the first timer 72 to start a countdown.

Please replace the paragraph beginning on page 89, line 12 with the following rewritten paragraph:

With reference to FIG. 14a, if four data blocks have been stored, all data blocks of the disk drives 5A to 5D have been stored in the buffer memories before the time margin T_{MARGIN} calculated based on the time $t_{4\text{th}}$ is consumed (that is, by the time $(t_{\text{ISSUE}} + t_0)$). Further, reading of the redundant data has also been completed. Therefore, the controller 7 is not required to issue a read termination command, and the procedure directly advances from step S44 to step S16. In step S16, the controller 7 issues a second READ-COMPLETED. In response to the second READ-COMPLETED, the selector 2 fetches the data blocks from the buffer memories 3A to 3D to assemble the data to be transmitted to the host device. The selector then transmits the assembled data through the host interface 1 to the host device. The first timer 72 stops the countdown, as required.

Please replace the paragraph beginning on page 90, line 9 with the following rewritten paragraph:

As described above, the disk array device of the third embodiment is different from that of the first embodiment in that an unread data block is not recovered immediately after four first READ-COMPLETED's arrive. In other words, the disk array device of the present embodiment waits until reading of the remaining data block has been completed within the time margin T_{MARGIN} after four first READ-COMPLETED's arrive. A recovery instruction is issued to the parity calculator 6 only after a lapse of the time margin T_{MARGIN} . When the remaining data block is read within the time margin T_{MARGIN} , four data blocks are stored in the buffer memories which allows the disk array device to transmit data to the host device without operating calculation of parity. Note that the time margin T_{MARGIN} is calculated, as described above with reference to FIG. 13a, based on the value t_0 which ensures that video being replayed at the host device is not interrupted. Furthermore, the time margin T_{MARGIN} indicates a time period within which reading of the remaining data block is expected to have been completed. Therefore, in most cases, four data blocks are stored in the buffer memories 3A to 3D within the time margin T_{MARGIN} . The present disk array seldom requires calculation of parity which requires a large amount of arithmetic operation, thereby minimizing the number of operations of calculation of parity.

Please replace the paragraph beginning on page 93, line 7 with the following rewritten paragraph:

The disk array device of the fourth embodiment performs a write operation as described in the first embodiment whenever transmission data from the host device arrives. To read data from the disk array device, the host device transmits a first read request specifying storage locations of the data to the disk array device.

Please replace the paragraph beginning on page 93, line 13 with the following rewritten paragraph:

In response to the first read request, the disk array device starts a read operation that is distinctive of the present embodiment which is now described in detail with reference to the drawings.

Please replace the paragraph beginning on page 94, line 7 with the following rewritten paragraph:

Further, the controller 7 regularly performs the procedure shown in a flow chart of FIG. 16. Since the flow chart of FIG. 16 partially includes the same steps as that of FIG. 12b, the same steps in FIG. 16 are provided with the same step numbers as those in FIG. 12b, and their descriptions are omitted herein.

Please replace the paragraph beginning on page 97, line 10 with the following rewritten paragraph:

Moreover, upon receiving the RECOVERY-COMPLETED, the controller 7 determines that four data blocks have been stored in the buffer memories and that the data requested from the host device can be now transmitted. The controller 7 then issues a second READ-COMPLETED (step S16). As a result, the data assembled by the selector 2 is transmitted through the host interface 1 to the host device.

Please replace the paragraph beginning on page 101, line 6 with the following rewritten paragraph:

Moreover, when the remaining data block arrives by the time the timer 74 is timed-out, the controller 7 does not issue any recovery instruction but issues a second READ-COMPLETED to assemble the data from the four data blocks and transmit the same to the host device. Therefore, the disk array device can minimize the number of operations of calculation of parity which requires a large amount of arithmetic operation.

Please replace the paragraph beginning on page 102, line 9 with the following rewritten paragraph:

In response to the first read request, the disk array device starts a read operation that is distinctive of the present embodiment, which is now described in detail with reference to a flow chart in FIG. 21. Since FIG. 21 partially includes the same steps as those in FIG. 2a, the same steps in FIG.

21 are provided with the same step numbers as those in FIG. 2a and their description is simplified herein.

Please replace the paragraph beginning on page 104, line 1 with the following rewritten paragraph:

When first READ-COMPLETE D's from the disk drives 5A to 5D and 5P arrive, the controller 7 performs an operation as shown in FIG. 2b. When any data block is failed to be read during this operation, the storage location of that data block is added to the faulty block table 75.

Please replace the paragraph beginning on page 105, line 1 with the following rewritten paragraph:

As described above, in the present disk array device, the redundant data is read only when the data blocks which have been failed to be read are to be read this time. Therefore, as described with reference to FIGS. 23a and 23b, the present disk array device can read a larger volume of data per unit of time. Furthermore, since the redundant data is read when there is a high possibility of failing to read the data blocks, the present disk array device can readily operate calculation of parity when the reading is actually failed, and transmit data to the host device as soon as possible.

Please replace the paragraph beginning on page 106, line 17 with the following rewritten paragraph:

As known, each of the disk drives 5A to 5D and 5P manages its own recording area by a sector unit of a predetermined size (512 bytes, in the present embodiment). A number called LBA is 20 assigned to each sector. LBA is an acronym for Logical Block Address. At initialization of the disk array device, part of the sectors in the recording areas of the disk drives are allocated for the alternate areas. The first table storage part 9 manages a first table 91 shown in FIG. 25 to manage such alternate areas. In FIG. 25, the LBA's specifying the allocated alternate areas are registered in the first table 91.

Please replace the paragraph beginning on page 107, line 23 with the following rewritten paragraph:

In response to the first read request, the present disk array device starts a read operation that is distinctive of the sixth embodiment which is now described with reference to FIG. 26. FIG. 26 shows a flow chart showing the procedure of the controller 7 after the first read request arrives. Since the flow chart of FIG. 26 partially includes the same steps as those of FIG. 12, the steps of FIG. 26 are provided with the same step numbers as those of FIG. 12 and their description is simplified herein.

Please replace the paragraph beginning on page 110, line 20 with the following rewritten paragraph:

When reading has been successfully completed, the disk drives 5A to 5D and 5P transmit the read data blocks and redundant data to the SCSI interfaces 4A to 4D and 4P. Further, each disk drive transmits an ACK, a read response indicating that reading has been successfully completed, to its corresponding SCSI interface. Upon receiving the ACK, each SCSI interface identifies which second read request the received ACK corresponds to, and stores the read data block or redundant data in the corresponding one of the buffer areas 3A_i to 3D_i and 3P (refer to FIG. 2) specified by the controller 7. Further, each SCSI interface transmits the received ACK to the controller 7.

Please replace the paragraph beginning on page 115, line 15 with the following rewritten paragraph:

When the procedure advances from step S82 to S88, a NAK has arrived at the controller 7. The controller 7 determines in step S88 whether the first timer 72 is active or not. When determining that the first timer 72 is not active, the procedure advances to step S811, which will be described later. On the other hand, when determining that the first timer 72 is active, the NAK indicates that reading of the remaining data block which had not yet been completed in step S814 has been eventually failed thereafter. The controller 7 realizes that countdown by the first timer 72 is no longer necessary, and stops the countdown (step S89). The controller 7 also realizes that reading of the remaining data block has failed and that the data block has to be recovered. The controller 7 thus

issues a recovery instruction to the parity calculator 6 for operating calculation of parity (step S810). The parity calculator 6 recovers the remaining unread data block, and stores the same in the buffer memory 3P. The parity calculator 6 then issues a RECOVERY-COMPLETED, a signal indicating that recovery of the data block has been successfully completed, to the controller 7. In response to the RECOVERY-COMPLETED, the controller 7 issues a READ-COMPLETED to the selector 2 (step S86). As a result, the data is transmitted to the host device.

Please replace the paragraph beginning on page 118, line 20 with the following rewritten paragraph:

Described next is operation of the reassignment part 8 with reference to FIGS. 29 to 34. As described above, the SCSI interfaces 4A to 4D and 4P are additionally provided with a new structure relating to the reassignment part 8. The new structure includes, as shown in FIG. 29, notifying parts 42A to 42D and 42P. When the SCSI interfaces 4A to 4D and 4P transmit second read requests to the disk drives 5A to 5D and 5P, respectively, each of the notifying parts 42A to 42D and 42P generates a transmission notification indicating the transmission of the second read request. The generated notifications are transmitted to the reassignment part 8. Each notification includes an ID uniquely specifying the transmitted second read request, and the LBA specified by the second read request. When the SCSI interfaces 4A to 4D and 4P receive a read response (ACK or NAK) from the disk drives 5A to 5D and 5P, respectively, each of the notifying parts 42A to 42D and 42P further generates a receive notification indicating the receiving of the read response. The generated receive notifications are transmitted to the reassignment part 8. Each receive notification includes an ID uniquely specifying the second read request corresponding to the received read response, and the LBA specified by the second read request. The reassignment part 8 can operate correctly, even if the LBA is not included in the receive notification.

Please replace the paragraph beginning on page 119, line 18 with the following rewritten paragraph:

Moreover, the reassignment part 8 includes, as shown in FIG. 29, a third timer 81 indicating the present time of day, a first list 82, and a second list 83, executing the procedure for reassignment shown in a flow chart of FIG. 30 whenever the reassignment part 8 receives a transmission notification. For a specific description, assume herein that the reassignment part 8 receives a transmission notification from the SCSI interface 4A. The received transmission notification includes the ID "b" and the LBA "a".

Please replace the paragraph beginning on page 126, line 5 with the following rewritten paragraph:

When the reassignment part 8 determines in step S103 that plural second read requests are kept, that is, plural first lists 82 are managed in the reassignment part 8, the procedure advances to step S104. Here, plural first lists 82 are managed for the disk drive 5A to be processed. Further, in step S108 or S1013 described later, the selected first list 82 is deleted. Therefore, at this time, as shown in FIG. 31 (a-3), the reassignment part 8 manages the first list 82 to be processed and the first list 82 created next (hereinafter referred to as "next first list 82") therein. The next first list 82 is shown as surrounded by a dotted line in FIG. 31 (a-3). Note that the next first list 82 does not include the process start time registered because it was created in step S94 of FIG. 30. To register the process start time, the reassignment part 8 first obtains the present time T_0 from the third timer 81, and registers the present time T_0 in the next first list 82 (step S104). The procedure then advances to step S105.

Please replace the paragraph beginning on page 127, line 7 with the following rewritten paragraph:

As described above, the second list 83 includes the fields for registering the LBA and the counter value N therein. The counter value N indicates how many times the LBA to be checked has successively satisfied $T_D > T_L$ in step S101. Therefore, if any second list 83 with the LBA to be checked registered therein is found in step S105, the LBA to be checked is determined to be possibly

defective also at a previous check. That is, the second read request for reading the data block or redundant data from the LBA to be checked has been transmitted successively at least twice (at previous time and this time) by now. Moreover, the reassignment part 8 has successively determined that the LBA to be checked satisfies $T_D > T_0$, twice in step S101 executed in response to each second read request. On the other hand, when any second list 83 with the LBA to be checked registered therein cannot be found, the LBA to be checked is determined for the first time to possibly be defective.

Please replace the paragraph beginning on page 129, line 20 with the following rewritten paragraph:

Upon receiving the transmission notification, the reassignment part 8 starts the procedure as shown in FIG. 30, first obtaining the present time T_0 from the third timer 81. The present time T_0 is used, as described above, as the time when the SCSI interface 4A transmits the second read request to the disk drive 5A. Here, assume that the transmission time of the second read request is t_2 . The reassignment part 8 extracts ID "c" and the LBA "a" from the received transmission notification (step S91). The reassignment part 8 then executes steps S92 and then S93, or steps S92 and then S94 to create a new first list 82 for the present second read request, and then ends the procedure of FIG. 30. Assuming that the present target disk drive (disk drive 5A) keeps only one second read request, the first list 82 includes the LBA "a", the ID "c", and the process start time " t_2 " registered therein (refer to FIG. 31 (a-4)).

Please replace the paragraph beginning on page 135, line 15 with the following rewritten paragraph:

On the other hand, when the reassignment part 8 finds the first list 82 to be deleted in step S111, $T_D > T_L$ has not been satisfied in step S101 of FIG. 32 by the time immediately before receiving the receive notification (that is, immediately before the present read response is returned thereto). Thus, the reassignment part 8 determines whether $T_D > T_L$ is satisfied or not at this time based on the information registered in the first list 82 to be deleted (step S112). When the delay time T_D exceeds the limit time T_L , the reassignment part 8 has to determine whether the alternate area has to be

assigned to the defective area, and the procedure therefore advances to steps S103 and thereafter as shown in FIG. 32, which are shown by "B" in the flow chart of FIG. 33.

Please replace the paragraph beginning on page 136, line 3 with the following rewritten paragraph:

On the other hand, when the delay time T_D does not exceed the limit time T_L , that means the reading of the disk drive 5A does not take a long time, and the LBA specified by "a" is not defective. Therefore, the reassignment part 8 determines whether the reassignment part 8 manages the second list 83 in which the same LBA as that in the first list 82 to be deleted is registered (step S113). When managing such second list 83, the reassignment part 8 deletes the second list 83 (step S114), and the procedure advances to step S115. Otherwise, the procedure directly advances from step S113 to step S115, wherein the reassignment part 8 deletes the first list 82 to be deleted.

Please replace the paragraph beginning on page 138, line 13 with the following rewritten paragraph:

On the other hand, when determining that the reading from the LBA has not yet been started, the reassignment part 8 transmits an ABORT_TAG message, one of the SCSI messages, to the disk drive 5 including the extracted LBA through the SCSI interface 4, terminating the execution of processing of the second read request corresponding to the found first list 82 (step S122). The SCSI interface 4 also transmits a NAK, indicating that the reading for the corresponding second read request has been failed, to the controller 7.

Please replace the paragraph beginning on page 138, line 22 with the following rewritten paragraph:

After step S122, the reassignment part 8 deletes the first list 82 found in step S121 (step S123).

Please replace the paragraph beginning on page 141, line 13 with the following rewritten paragraph:

When determining in step S131 that reading of the four data blocks has not previously failed, the controller 7 determines that there is a low possibility of failing to read the four data blocks this time, and issues a set of second read requests to read the parity group (step S132). In step S132, however, the second read requests are issued only to the four disk drives storing the data blocks, and not to the remaining disk drive storing the redundant data.

Please replace the paragraph beginning on page 141, line 21 with the following rewritten paragraph:

When determining in step S131 that reading of the four data blocks has previously failed, the controller 7 determines that there is a high possibility of failing to read the four data blocks this time, and issues a set of second read requests to read the parity group (step S133). In step S133, however, the second read requests are issued to the four disk drives storing the data blocks as well as the remaining disk drive storing the redundant data.

Please replace the paragraph beginning on page 142, line 4 with the following rewritten paragraph:

The second read requests issued in step S132 are processed by the four disk drives storing the data blocks of the same parity group, while those issued in step S133 are processed by the five disk drives storing the data blocks and redundant data of the same parity group. In either case, each of the four or five disk drives generates a read response indicating reading has succeeded or failed. The four or five disk drives transmit the generated read responses through the SCSI interfaces connected thereto to the controller 7. The controller 7 executes the procedure shown in FIG. 37 whenever the read response arrives. The flow chart of FIG. 37 includes the same steps as those in the flow chart of FIG. 28, and further includes step S141. Therefore, the steps in FIG. 37 are provided with the same step numbers as those in FIG. 28 and their description is omitted herein.

Please replace the paragraph beginning on page 142, line 18 with the following rewritten paragraph:

When determining that a NAK has arrived (step S82), the controller 7 extracts the LBA from the NAK. The LBA included in the NAK indicates the storage location of the data block or redundant data which has failed to be read. The controller 7 registers the LBA extracted from the NAK in the faulty block table 75 (step S141). Note that step S141 may be executed at any time as long as after it is determined in step S82 that the present read response is a NAK. That is, the execution timing of step S141 is not restricted to the timing immediately after determined in step S82 that the present read response is a NAK.

Please replace the paragraph beginning on page 143, line 3 with the following rewritten paragraph:

The reassignment part 8 executes the procedure described above in the sixth embodiment. Description of this procedure is therefore omitted herein. The important point here is that, when the reassignment ends, the reassignment part 8 transmits a REASSIGN-COMPLETED notification, indicating the reassignment has ended, to the controller 7. This REASSIGN-COMPLETED notification includes the LBA indicative of the storage location that is determined to be defective by the reassignment part 8. Since it takes much time to read from the defective area, the LBA indicative of such defective storage area is also written in the faulty block table 75.

Please replace the paragraph beginning on page 143, line 14 with the following rewritten paragraph:

When receiving the REASSIGN-COMPLETED notification, the controller 7 executes the procedure shown in FIG. 38. First, upon receiving REASSIGN-COMPLETED notification, the controller 7 determines that the reassignment part 8 has executed reassignment (step S151), and the procedure advances to step S152. In step S152, the controller 7 extracts the LBA from the REASSIGN-COMPLETED notification. The controller 7 then accesses to the faulty block table 75, and deletes the LBA matching the one extracted from the REASSIGN-COMPLETED notification from the faulty block table 75, thereby updating the faulty block table 75 (step S152).

Please replace the paragraph beginning on page 144, line 17 with the following rewritten paragraph:

In the above described first to seventh embodiments, the disk array device includes five disk drive. The number of disk drives, however, may be changed according to design requirements of the disk array device such as the data length and the data block length, and therefore is not restricted to five. Note that "m" in the claims corresponds to the number of disk drives included in the disk array device.

Please replace the paragraph beginning on page 145, line 8 with the following rewritten paragraph:

As described in the Background Art section, the disk array device executes reconstruction processing in some cases. In an eighth embodiment of the present invention, reconstruction is to recover the data block or redundant data in a faulty disk drive and rewrite the recovered data block or redundant data in a disk drive (another disk drive or a recording area without a defect in the faulty disk drive). Furthermore, the disk array device has to transmit video data so that the video being replayed at the host device is not interrupted. To prevent this interruption of video, when a read request for video data arrives, the disk array device has to process the read request in real time to transmit the video data. The eighth embodiment realizes a disk array device capable of transmitting video data without interruption and executing reconstruction.

Please replace the paragraph beginning on page 147, line 1 with the following rewritten paragraph:

Described next is the detailed structure of the buffer managing part 37 with reference to FIG. 41. In FIG. 41, the buffer managing part 37 manages buffer memories 37A to 37D, 37P, and 37R. The buffer memory 37A is divided into a plurality of buffer areas 37A₁, 37A₂... each buffer area having a capacity of storing a data block or redundant data, which will be described below. Further, an identifier (normally, top address of each buffer area) is assigned to each buffer area to uniquely identify each buffer area. The identifier of each buffer area is hereinafter referred to as a pointer.

Each of the other buffer memories 37B to 37D, 37P, and 37R is also divided into a plurality of buffer areas. A pointer is also assigned to each buffer area, like the buffer area 37A₁.

Please replace the paragraph beginning on page 147, line 14 with the following rewritten paragraph:

Referring back to FIG. 40, the disk group of the disk drives 41A to 41D and 41P is now described. Since the architecture of the present disk array device is based on the combination of RAID-3 and RAID-4, the data blocks and redundant data of the same parity group are distributed across the disk drives 41A to 41D and 41P, which form one disk group. Here, the parity group is, as described in the Background Art section, a set of data blocks and redundant data generated based on one piece of data transmitted from the host device. The disk group is a set of plurality of disk drives into which the data blocks and redundant data of the same parity group are written. In the present embodiment, the disk group of the disk drives 41A to 41D and 41P is hereinafter referred to as a disk group "A". Further, a plurality of LUN's (Logical Unit Number) are assigned to each disk group. The plurality of LUN's are different for each disk group, and the LUN's in one disk group are also different each other. Such LUN's are used for specifying a disk group to be accessed and the level of priority of an access request. In the present embodiment, "non-priority" and "priority" are previously defined as the level of priority of an access request. Two LUN's "0" and "1" are assigned to the disk group A. The LUN "0" represents that the access request is given "non-priority", while the LUN "1" represents the access request is given "priority".

Please replace the paragraph beginning on page 171, line 14 with the following rewritten paragraph:

In these two types of reconstruction, a table storage part 39 shown in FIG. 49 is used. The table storage part 39, as shown in FIG. 49, stores managing tables 39A to 39D and 39P for the disk drives 41A to 41D and 41P (the disk group "A"). LBA statuses assigned to each entire recording area of the disk drives 39A to 39D and 39P are stored in the managing tables 39A to 39D and 39P, respectively. For example, the LBA status is set in each corresponding section in the managing table 39A.

Please replace the paragraph beginning on page 172, line 4 with the following rewritten paragraph:

When detecting that one of the disk drives 41A to 41D and 41P failed, the SCSI interfaces 36A to 36D and 36P first notifies the controller 33 that the disk drive is defective. Here, the faulty disk drive is detected when a notification of the faulty disk drive is received or when a response from the disk drives 41A to 41D to 41P does not return to the SCSI interfaces 36A to 36D and 36P within a predetermined time.

Please replace the paragraph beginning on page 179, line 20 with the following rewritten paragraph:

When determining in step S212 that the predetermined time T has elapsed, the controller 33 executes the second reconstruction processing for the LBA to be processed (step S212). FIG. 54 is a flow chart showing the detailed procedure in step S212. FIG. 54 is different from FIG. 52 only in that steps S202 and S207 are replaced with steps S221 and S222. Therefore, in FIG. 54, the steps corresponding to the steps in FIG. 52 are provided with the same step numbers as those in FIG. 52 and their description is simplified herein.

Please replace the paragraph beginning on page 181, line 3 with the following rewritten paragraph:

The request selection unit 35A executes the same processing as in step S169 (step S208). Therefore, the present sixth write request is dequeued from the priority queue 342A by the request selection unit 35A and transmitted to the SCSI interface 36A. The SCSI interface 36A processes the received sixth write request, and the disk drive 41A writes the recovered data in the LBA to be processed (step S209). In this way, enqueued to the priority queue 342A, the present sixth write request is processed by the disk drive 41A with priority. When completing the write operation, the disk drive 41A transmits a WRITE-COMPLETED, a signal indicating that writing has been completed, to the controller 33 through the SCSI interface 36A.

Please replace the paragraph beginning on page 181, line 15 with the following rewritten paragraph:

The controller 33 then executes steps S2010 and S2011, bringing the processing of step S194 to an end. Furthermore, the controller 33 executes the loop of steps S192 to S196 until all of the LBAs are subjected to the processing of step S194.

Please replace the paragraph beginning on page 181, line 19 with the following rewritten paragraph:

According to the second reconstruction, the requests for reconstruction (seventh read request and sixth write request) are enqueued to the priority queues. This can shorten the time the request waits to be processed in the queue managing part 34, thereby ensuring the time when the data is fully reconstructed. Furthermore, the array controller 21 enqueues each request and controls the second reconstruction processing for each disk drive, thereby effectively performing the second reconstruction processing.

Please replace the paragraph beginning on page 183, line 24 with the following rewritten paragraph:

As described above, when the host device requests access to the LBA "reconstruction-required" in the newly-created managing table, the disk array device writes the data recovered with calculation of parity in the LBA. The write request for this writing is enqueued to the non-priority queue. Therefore, the recovered data is written in the disk array 22 with lower priority together with the access request from the host device. As described above, the LBA "reconstruction-required" is subjected to the first or second reconstruction processing. However, the first and second reconstruction processings are executed in parallel, decreasing the number of LBA "reconstruction-required" in either processing. This shortens the time required for the first or second reconstruction processing. Furthermore, since the seventh write request is enqueued to the non-priority queue, it can be ensured that writing of the recovered data does not affect other processing with higher priority to be executed by the disk array device.

Please replace the paragraph beginning on page 184, line 16 with the following rewritten paragraph:

When the host device requests access to the LBA "reconstruction-required" for writing the data, the controller 33 changes the status of the LBA to "normal" when the disk array device completes writing. Therefore, the disk array device is not required to execute unnecessary reconstruction processing, and the processing time in the disk array device can be shortened.

Please replace the paragraph beginning on page 194, line 6 with the following rewritten paragraph:

Described next is the processing in step S251 in detail, taking the ID "b" for example. In the first list 751 (refer to FIG. 60 (a-2)), the I/O request SSR 1 is specified by the ID "b", and its delay time T_{D1} therefore can be given by $T_0 - T_{T1}$. When $T_{D1} > T_L$ is satisfied, the procedure advances to step S252. When not satisfied, the reassignment part 75 executes the processing in step S251 again to find the ID for reassignment. Note again that, in step S251, the reassignment part 75 does not monitor for the I/O request SSR in which the process start time has not yet been registered.

Please replace the paragraph beginning on page 194, line 16 with the following rewritten paragraph:

When determining in step S251 that $T_{D1} > T_L$ is satisfied for the ID "b", the reassignment part 75 instructs the disk interface control part 74 to terminate execution of the I/O request SSR 1 specified by the ID "b" (step S252). In response to this instruction, the disk interface 74 transmits a ABORT_TAG message, which is one of the SCSI messages, to terminate execution of the I/O request SSR 1. The disk interface 74 then notifies the read/write controller 73 that the processing of the I/O request SSR 1 has failed. In response, the read/write controller 73 executes the processing, which will be described later.

Please replace the paragraph beginning on page 198, line 2 with the following rewritten paragraph:

When the response RES 1 indicates that writing has failed, the read/write controller 73 generates an I/O request SSR 1' including the same information as the I/O request SSR 1 to retry to register the sub-segment in the successive LBA area "a", and then transmits the same to the disk drive 62. When the response RES 1 indicates that reading has failed, the read/write controller 73 recovers the unread sub-segment or retries to register the sub-segment as described above by using parity and other sub-segments according to the RAID architecture.

Please replace the paragraph beginning on page 202, line 4 with the following rewritten paragraph:

Upon receiving the affirmative response ACK 1, the read/write controller 73 instructs the disk drive 62 subjected to reassignment to write the sub-segment when the I/O request SSR 1' requests write operation. When the I/O request SSR1' requests read operation, the read/write controller 73 recovers the sub-segment lost by reassignment using parity and other sub-segments according to the RAID architecture, and then transmits the recovered sub-segment to the host device 81 through the host interface 72 and also instructs the disk drive 62 through the disk interface 74 to write the recovered sub-segment. Thus, the recorded data in the disk drive 62 can maintain consistency before and after reassignment.

Please replace the paragraph beginning on page 203, line 2 with the following rewritten paragraph:

The disk interface 74 generates a signal "receive notification" whenever it receives the response RES to the I/O request SSR, and transmits the same to the reassignment part 75. This receive notification includes the ID and successive LBA area of the I/O request on which the received response RES is based. The reassignment part 75 executes the flow chart of FIG. 59 whenever it receives a receive notification. Now, assume herein that the disk interface 74 generates the receive notification upon receiving the response RES 1 and transmits the same to the reassignment part 75. The response RES 1 includes, as evident from above, the ID "b", the successive LBA information "a"

and the information on whether read or write. Note that the information on whether read or write is not required for the reassignment part 75. Therefore, the receive notification only includes the ID "b" and the LBA "a".

Please replace the paragraph beginning on page 219, line 19 with the following rewritten paragraph:

In step S2715, when determining that the disk drive 1002 is unsuitable for recording data, the reassignment part 1103 ceases to use the disk drive 1002 for data recording, and determines to use the spare disk drive 1003. In response to this determination, the disk controller 1101 controls the disk group 1001 to copy the data (sub-segment, parity, data recorded in the system area) recorded in the disk drive 1002 to the spare disk drive 1003 (step S2716). After this copy control ends, the disk controller 1101 updates the address information 11110 to provide consistency in the original LBA and the current LBA. Thus, even if receiving the I/O request SR specifying the original LBA from the host device 81, the read/write controller 1102 can fetch the current LBA of the sub-segment from the address conversion part 1107. In other words, the disk controller 1101 can correctly recognize the spare disk drive 1003 as the disk drive for recording data. Therefore, the host device 81 is not required to recognize the replacement of the disk drive 1002 with the spare disk drive 1003 in the disk group 1001.